

Administering vRealize Automation

21 December 2020
vRealize Automation 8.0

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

- 1 Administering vRealize Automation** 4
- 2 Administering users** 5
 - [How do I enable Active Directory groups in vRealize Automation for projects](#) 6
 - [How do I remove users in vRealize Automation](#) 7
 - [How do I edit user roles in vRealize Automation](#) 7
 - [How do I edit group role assignments in vRealize Automation](#) 8
- 3 Maintaining your appliance** 9
 - [Starting and stopping vRealize Automation](#) 9
 - [How do I enable time synchronization](#) 11
 - [How do I deactivate time synchronization](#) 12
 - [How do I reset the root password](#) 12
- 4 Working with logs** 15
 - [How do I work with logs and log bundles](#) 15
 - [How do I configure log forwarding to vRealize Log Insight](#) 17
- 5 Participating in the Customer Experience Improvement Program** 21
 - [How do I join or leave the program](#) 21
 - [How do I configure the data collection time for the program](#) 22

Administering vRealize Automation

1

While most vRealize Automation administration tasks are completed from VMware vRealize Suite Lifecycle Manager, this guide describes some important user and system management tasks that you can complete from within vRealize Automation.

For more information about working with vRealize Suite Lifecycle Manager, see [vRealize Suite Lifecycle Manager installation, upgrade, and management](#).

While some vRealize Automation administration tasks are completed from within vRealize Automation, others require the use of related products such as vRealize Suite Lifecycle Manager and Workspace ONE Access. Users should familiarize themselves with these products and their functionality before completing applicable tasks.

For example, for information about backup, restore, and disaster recovery, see the **Backup and Restore, and Disaster Recovery > 2019** section of [vRealize Suite product documentation](#).

Note Disaster recovery is not supported by vRealize Automation 8.0.0. To use vRealize Automation in disaster recovery scenarios, upgrade to vRealize Automation 8.0.1 or later.

Administering Users and Groups in vRealize Automation

2

vRealize Automation uses VMware Workspace ONE Access, the VMware supplied identity management application to import and manage users and groups. After users and groups are imported or created, you can manage role assignments using the Identity & Access Management page.

vRealize Automation is installed using VMware Lifecycle Manager (vRSLCM or LCM). When installing vRealize Automation you must import an existing Workspace ONE Access instance, or deploy a new one to support identity management. These two scenarios define your management options.

- If you deploy a new Workspace ONE Access instance, you can manage users and groups via LCM. During installation, you can set up an Active Directory connection using Workspace ONE Access. Alternatively, you can view and edit some aspects of users and groups within vRealize Automation using the Identity & Access Management page as described herein.
- If you use an existing Workspace ONE Access instance, you import it for use with vRealize Automation via LCM during installation. In this case, you can continue to use Workspace ONE Access to manage users and groups, or you can use the management functions in LCM.

vRealize Automation users must be assigned roles. Roles define access to features within the application. When vRealize Automation is installed with a Workspace ONE Access instance, a default organization is created and the installer is assigned the Organization Owner role. All other vRealize Automation roles are assigned by the Organization Owner.

There are three types of roles in vRealize Automation: organization roles, service roles, and project roles. For vRealize Automation Cloud Assembly, Service Broker and Code Stream, typically, user level roles can use resources, while admin level roles are required to create and configure resources. Organizational roles define permissions within the tenant; organizational owners have admin level permissions while organizational members have user level permissions. Organization owners can add and manage other users.

Organization Roles	Service Roles
■ Organization Owner	■ Cloud Assembly Administrator
■ Organization Member	■ Cloud Assembly User
	■ Service Broker Administrator
	■ Service Broker User
	■ Code Stream Administrator
	■ Code Stream User
	■ Code Stream Viewer

In addition, there are two main project level roles not shown in the table: Project Administrator, and Project User. These roles are assigned ad hoc on a per project basis with Cloud Assembly. These roles are somewhat fluid. The same user can be an administrator on one project and a user on another project.

For more information about working with LCM and Workspace ONE Access, see [User Management with VMware Identity Manager](#).

This chapter includes the following topics:

- [How do I enable Active Directory groups in vRealize Automation for projects](#)
- [How do I remove users in vRealize Automation](#)
- [How do I edit user roles in vRealize Automation](#)
- [How do I edit group role assignments in vRealize Automation](#)

How do I enable Active Directory groups in vRealize Automation for projects

If a group is not available on the Add Groups page when you are adding users to projects, check the Identity & Access Management page and add the group if it is available. If the group is not listed on the Identity & Access Management page in vRealize Automation, the group may not be synchronized in your Workspace One Access instance. You can verify that it has been synchronized and then use this procedure to add the group as shown herein.

To add members of an Active Directory group to a project, you must ensure that the group is synchronized with your Workspace One Access instance and that the group is added to the organization.

Prerequisites

If the groups are not synchronized, they are not available when you try to add them to a project. Verify that you synchronized your Active Directory groups with your Lifecycle Manager instance.

Procedure

- 1 Log in to vRealize Automation as a user from the same Active Directory domain that you are adding. For example, @mycompany.com

- 2 In Cloud Assembly, click Identity & Access Management in the header right navigation.
- 3 Click **Enterprise Groups**, and then click **Assign Roles**.
- 4 Use the search function to find the group that you are adding and select it.
- 5 Assign an organization role.

At a minimum, the group must have an Organization Member role. See [What are the Cloud Assembly user roles](#) for more information.

- 6 Click **Add Service Access**, add one or more services, and select a role for each.
- 7 Click **Assign**.

Results

You can now add the Active Directory group to a project.

How do I remove users in vRealize Automation

You can remove users as needed in vRealize Automation.

All users are listed by default and you cannot add users with the Identity and Access Management page. You can delete users.

Procedure

- 1 Select the Active Users tab on the Identity & Access Management page.
- 2 Locate and select the users that you want to delete.
- 3 Click **Remove Users**.

Results

The selected users are removed.

How do I edit user roles in vRealize Automation

You can edit roles assigned to Workspace One Access users that have been imported into vRealize Automation.

Prerequisites

Procedure

- 1 In Cloud Assembly, click Identity & Access Management in the header right navigation.
- 2 Select the desired user on the Active Users tab and click **Edit Roles**.

- 3 You can edit the organization and service roles for the user.
 - Select the drop down beside the Assign Organization Roles heading to change the user's relationship with the organization.
 - Click Add Service Access to add new service roles for the user.
 - To remove user roles, click the X beside the applicable service.
- 4 Click **Save**.

Results

The user role assignment is updated as specified.

How do I edit group role assignments in vRealize Automation

You can edit role assignments for groups in vRealize Automation

Prerequisites

Users and groups have been imported from a valid vIDM instance that is associated with your vRealize Automation deployment.

Procedure

- 1 In Cloud Assembly, click Identity & Access Management in the header right navigation.
- 2 Select the Enterprise Groups tab.
- 3 Type the name of the group for which you want to edit role assignments in the search field.
- 4 Edit the role assignments for the selected group. You have two options.
 - Assign Organization Roles
 - Assign Service Roles
- 5 Click **Assign**.

Results

Role assignments are updated as specified.

Maintaining your vRealize Automation appliance

3

As a system administrator, you might need to perform various tasks to ensure the proper functioning of your installed vRealize Automation application.

If you are just getting started with vRealize Automation, these are not required tasks. Knowing how to perform these tasks is useful if you need to resolve performance or product behavior issues.

This chapter includes the following topics:

- [Starting and stopping vRealize Automation](#)
- [How do I enable time synchronization of vRealize Automation](#)
- [How do I deactivate time synchronization](#)
- [How do I reset the root password for vRealize Automation](#)

Starting and stopping vRealize Automation

Observe the proper procedures when starting or shutting down vRealize Automation.

Shut down vRealize Automation

To preserve data integrity, shut down the vRealize Automation services before powering off the virtual appliances.

Note Avoid using `vraccli reset vidm` commands if at all possible. This command resets all configuration of Workspace One Access and breaks the association between users and provisioned resources.

- 1 Log in to the console of any vRealize Automation appliance using either SSH or VMRC.

- 2 To shut down the vRealize Automation services on all cluster nodes, Run the following set of commands.

Note If you copy any of these commands to run and they fail, paste them into notepad first, and then copy them again before running them. This procedure strips out any hidden characters and other artifacts that might exist in the documentation source.

```
/opt/scripts/svc-stop.sh  
sleep 120  
/opt/scripts/deploy.sh --onlyClean
```

- 3 Shut down the vRealize Automation appliances.

Your vRealize Automation deployment is now shut down.

Start vRealize Automation

Following an unplanned shutdown, a controlled shutdown, or a recovery procedure, you must restart vRealize Automation components in a specific order. vRLCM is a non-critical component, so you can start it at any time. VMware Workspace ONE Access, formerly VMware Identity Management, components must be started before you start vRealize Automation.

Note Verify that applicable load balancers are running before starting vRealize Automation components.

- 1 Power on all vRealize Automation appliances and wait for them to start.
- 2 Log into the console for any appliance using SSH or VMRC and run the following command to restore the services on all nodes.

```
/opt/scripts/deploy.sh
```

- 3 Verify that all services are up and running with the following command.

```
kubect1 get pods --all-namespaces
```

Note You should see three instances of every service, with a status of either Running or Completed.

When all services are listed as Running or Completed, vRealize Automation is ready to use.

Restart vRealize Automation

You can restart all vRealize Automation services centrally from any of the appliances in your cluster. Follow the preceding instructions to shut down vRealize Automation, and then use the instructions to start vRealize Automation. Before restarting vRealize Automation, verify that all applicable load balancer and VMware Workspace ONE Access components are running.

When all services are listed as Running or Completed, then vRealize Automation is ready to use.

Run the following command to verify that all services are running:

```
kubectl -n prelude get pods
```

How do I enable time synchronization of vRealize Automation

You can enable time synchronization on your vRealize Automation deployment by using the vRealize Automation Appliance command line.

You can configure time synchronization for your standalone or clustered vRealize Automation deployment by using the Network Time Protocol (NTP) networking protocol. vRealize Automation supports two mutually exclusive NTP configurations:

NTP configuration	Description
ESXi	<p>You can use this configuration when the ESXi server hosting the vRealize Automation Appliance is synchronized with an NTP server. If you are using a clustered deployment, all ESXi hosts must be synchronized with an NTP server.</p> <p>Note You can experience clock drift if your vRealize Automation deployment is migrated to a ESXi host that is not synchronized to an NTP server.</p> <p>For more information on configuring NTP for ESXi, see KB article 57147 Configuring Network Time Protocol (NTP) on an ESXi host using the vSphere Web Client.</p>
systemd	<p>This configuration uses the systemd-timesyncd daemon to synchronize the clocks in your vRealize Automation deployment.</p> <p>Note By default, the systemd-timesyncd daemon is enabled, but configured with no NTP servers. If the vRealize Automation Appliance uses a dynamic IP configuration, the appliance can use any NTP servers received by the DHCP protocol.</p>

Procedure

- 1 Log in to the vRealize Automation Appliance command line as **root**.
- 2 Enable NTP with ESXi.
 - a Run the `vraccli ntp esxi --enable` command.
 - b Run the `vraccli ntp apply` command.

The ESXi NTP configuration is applied to the vRealize Automation deployment.

3 Enable NTP with systemd.

- a Run the `vraccli ntp systemd --set FQDN_or_IP_of_systemd_server` command.

Note You can add multiple systemd NTP servers by separating their network addresses with a comma.

- b Run the `vraccli ntp apply` command.

The systemd NTP configuration is applied to the vRealize Automation deployment.

4 (Optional) To confirm the status of the NTP configuration, run the `vraccli ntp status` command.

The NTP configuration can fail if there is a time difference of more than 10 minutes between the NTP server and the vRealize Automation deployment. To resolve this issue, reboot the vRealize Automation Appliance that is synchronized with the NTP server.

How do I deactivate time synchronization

You can deactivate the Network Time Protocol (NTP) time synchronization on your vRealize Automation deployment with the vRealize Automation Appliance command line.

Prerequisites

Verify that you have configured time synchronization with ESXi or systemd. See [How do I enable time synchronization of vRealize Automation](#).

Procedure

1 Log in to the vRealize Automation Appliance command line as **root**.**2** Deactivate a ESXi NTP configuration.

- a Run the `vraccli ntp esxi --disable` command.
- b Run the `vraccli ntp apply` command.

The ESXi NTP configuration is deactivated.

3 Deactivate a systemd NTP configuration.

- a Run the `vraccli ntp systemd --disable FQDN_or_IP_of_systemd_server` command.
- b Run the `vraccli ntp apply` command.

The systemd NTP configuration is deactivated.

4 (Optional) To confirm the status of the NTP configuration, run the `vraccli ntp status` command.

How do I reset the root password for vRealize Automation

You can reset a lost or forgotten vRealize Automation root password.

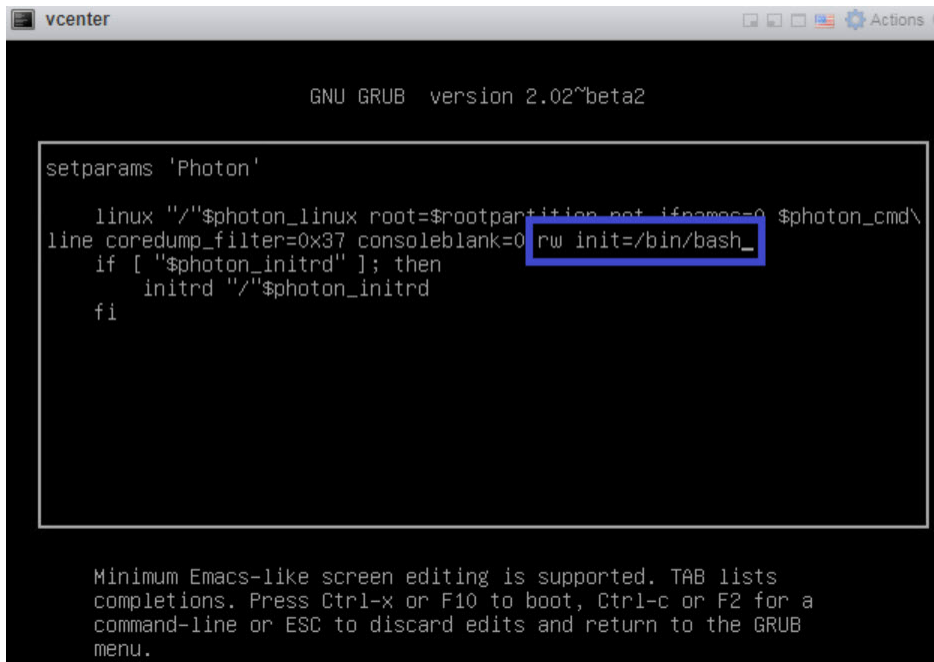
In this procedure, you use a command line window on the host vCenter appliance to reset your organization's vRealize Automation root password.

Prerequisites

This process is for vRealize Automation administrators and requires the credentials needed to access the host vCenter appliance.

Procedure

- 1 Shut down and start up vRealize Automation by using the procedure described in [Starting and stopping vRealize Automation](#).
- 2 When the Photon operating system command line window appears, enter `e` and press the **Enter** key to open the GNU GRUB boot menu editor.
- 3 In the GNU GRUB editor, enter `rw init=/bin/bash` at the end of the line that begins with `linux` `"/" $photon_linux root=rootpartition` as shown below:



```

GNU GRUB version 2.02~beta2

setparams 'Photon'

linux "/"$photon_linux root=$rootpartition root.ifnames=0 $photon_cmd\
line coredump_filter=0x37 consoleblank=0 rw init=/bin/bash_
if [ "$photon_initrd" ]; then
    initrd "/"$photon_initrd
fi

Minimum Emacs-like screen editing is supported. TAB lists
completions. Press Ctrl-x or F10 to boot, Ctrl-c or F2 for a
command-line or ESC to discard edits and return to the GRUB
menu.

```

- 4 Click the **F10** key to push your change and restart vRealize Automation.
- 5 Wait for vRealize Automation to restart.
- 6 At the root `[/]#` prompt, enter `passwd` and press the **Enter** key.
- 7 At the New password: prompt, enter your new password and press the **Enter** key.
- 8 At the Retype new password: prompt, reenter your new password and press the **Enter** key.

- 9 At the root `[/]#` prompt, enter `reboot -f` and press the **Enter** key to complete the root password reset process.

```
root [ / ]# passwd
New password:
Retype new password:
passwd: password updated successfully
root [ / ]# reboot -f_
```

What to do next

As a vRealize Automation administrator, you can now log in to vRealize Automation with the new root password.

Working with logs in vRealize Automation

4

You can use the supplied `vracli` command line utility to create and use logs in vRealize Automation.

You can use logs directly in vRealize Automation or you can instead forward all logs to vRealize Log Insight.

This chapter includes the following topics:

- [How do I work with logs and log bundles in vRealize Automation](#)
- [How do I configure log forwarding to vRealize Log Insight](#)

How do I work with logs and log bundles in vRealize Automation

You can create and use vRealize Automation logs and log bundles in vRealize Automation.

Alternatively, you can automatically forward logs to vRealize Log Insight. For information about how to forward logs to vRealize Log Insight, see [How do I configure log forwarding to vRealize Log Insight](#).

Information about how to use the `vracli` command line utility is available by using the `--help` argument in the `vracli` command line. For example: `vracli log-bundle --help`.

Log bundle commands

You can create a simple log bundle or an aggregated (cold storage) log of all services. While both log bundles contains all the logs for your services, the cold storage bundle contains a copy of an aggregated stream of back-versions of the service logs, which can supply additional troubleshooting value. The cold storage agent constantly aggregates logs from the services and stores them on the local file system. A simple log bundle is typically all that is needed for troubleshooting.

You can also change the default timeout value for collecting logs from each node.

In a clustered environment, you only need to run the `vracli log-bundle` command on one node.

- Display the log bundle command help:

```
vracli log-bundle --help
```

- Create a simple log bundle.

```
vraccli log-bundle
```

- Create a cold storage log bundle:

```
vraccli log-bundle --include-cold-storage
```

- Change the timeout value for collecting logs from each node. For example, if your environment contains large log files, slow networking, high CPU usage, and so on you might need to set the timeout to greater than the 1000 second default value.

```
vraccli log-bundle --collector-timeout $CUSTOM_TIMEOUT_IN_SECONDS
```

Log bundle structures

vRealize Automation services are containerized in Kubernetes pods. The generated log bundle is a `tar.xz` archive that uses a `log-bundle-{{TIMESTAMP}}.tar.xz` name format, where `TIMESTAMP` is an epoch timestamp in seconds. A normal log bundle contains logs from all the nodes in the environment. If the log bundle cannot be generated for whatever reason, a fallback bundle is created instead. The fallback bundle contains logs for the current node only. There are slight differences in the structure of the 2 log bundles types.

- Normal log bundles

Normal log bundles are organized into the following categories:

- Host logs and configuration

The configuration for each host and its host-specific logs are collected in one directory per cluster node (host). The directory name matches the node host name. The directory contents match the host file system. The number of directories matches the number of cluster nodes.

Cold storage logs are located in a structured JSON log as `/hostname/services-logs/all/aggregated.log`.

- Pod logs

Services are containerized in Kubernetes pods. Service logs are located in the `pods` directory, which contains a single directory per namespace with a file name that matches the namespace name. There is typically one instance of each pod per cluster node. The pod directory contains a log file for each of its container applications.

For example, vRealize Orchestrator Control Center logs reside in a `vco-controlcenter-app.log` file in each of the `/pods/prelude/vco-app-hash/` directories.

- Environment file

The environment file contains information about the current resource usage per nodes and per pods. It also contains cluster information and descriptions for all the available Kubernetes entities.

- Fallback log bundles

If you receive an error messages while waiting for the `vracli` command to finish, a fallback bundle is generated. If you receive this error, you should run the `vracli log-bundle` command on each host or node in the cluster to collect as much information as possible.

- Fallback container logs

Fallback logs are located in the `/fallback-containers` directory. You can identify which container in which pod generated the logs by examining the log file name:

pod-name-some-hash-container-name-other-hash.log

- Fallback cold storage

If you are collecting cold storage logs with the bundle, the fallback logs from the current host are located in the `/fallback-cold-storage` directory.

How do I configure log forwarding to vRealize Log Insight

You can forward logs from vRealize Automation to vRealize Log Insight to take advantage of more robust log analysis and report generation.

vRealize Automation is bundled with a [fluentd-based](#) logging agent. The agent collects and stores logs so that they can be included in a log bundle and examined later. You can configure the agent to forward a copy of the logs to a vRealize Log Insight server by using the vRealize Log Insight API. The supplied API allows other programs to communicate with vRealize Log Insight.

For more information about vRealize Log Insight, including documentation for the vRealize Log Insight API, see [vRealize Log Insight documentation](#) and also the `/api/v1/events/ingest/{agentId}` page.

Configure the logging agent to automatically and continuously forward vRealize Automation logs to vRealize Log Insight by using the supplied `vracli` command line utility.

Information about how to use the `vracli` command line utility is available by using the `--help` argument in the `vracli` command line. For example: `vracli vrli --help`.

Check existing configuration of vRealize Log Insight

Command

```
vracli vrli
```

Arguments

There are no command line arguments.

Output

The current configuration for vRealize Log Insight integration is output in JSON format.

Exit codes

The following exit codes are possible:

- 0 - Integration with vRealize Log Insight is configured.

- 1 - An exception occurred as part of command execution. Examine the error message for details.
- 61 (ENODATA) - Integration with vRealize Log Insight is not configured. Examine the error message for details.

Example – check integration configuration

```
$ vracli vrli
No vRLI integration configured

$ vracli vrli
{
  "agentId": "0",
  "environment": "prod",
  "host": "my-vrli.local",
  "port": 443,
  "scheme": "https",
  "sslVerify": false
}
```

Note You can set a different host scheme (the default is https) and port (the default is 443) to use for sending the logs, as shown in the following samples:

```
vracli vrli set some-host
vracli vrli set some-host:9543
vracli vrli set http://some-host:9543
```

Port 9543 is used by the vRealize Log Insight ingestion API as described in the *Administering vRealize Log Insight* topic *Ports and External Interfaces* in the [vRealize Log Insight documentation](#).

Configure or update integration of vRealize Log Insight

Command

```
vracli vrli set [options] FQDN_OR_URL
```

Arguments

The following command line arguments are available:

- FQDN_OR_URL - the FQDN or IP address of the vRealize Log Insight server that is to be used to post logs by using the vRealize Log Insight API configuration. Port 443 and an HTTPS scheme are used by default. If any of these settings must be changed, you can use a URL instead.
- options
 - `--agent-id SOME_ID` - Set the ID of the logging agent for this appliance. The default value is 0. Use to identify the logging agent for logs that are posted to vRealize Log Insight by using the vRealize Log Insight API configuration.

- `--environment ENV` - set an identifier for the current environment. It will be available in vRealize Log Insight logs as a tag for each log line event. The default value is `prod`.
- `--ca-file /path/to/server-ca.crt` - Specify a file that contains the certificate authority (CA) certificate that was used to sign the vRealize Log Insight server certificate. Force the logging agent to trust the specified CA and enable it to verify the certificate of the vRealize Log Insight server. The file can contain a whole certificate chain if needed to verify the certificate. In case of a self-signed certificate, pass the certificate itself.
- `--ca-cert CA_CERT` - Specify a file in the same manner as `--ca-file` but pass the certificate (chain) inline as a string.
- `--insecure` - Deactivate SSL verification of the server certificate. Force the logging agent to accept any SSL certificate when posting logs.

Output

No output is expected.

Exit codes

The following exit codes are possible:

- 0 - The configuration was updated.
- 1 - An exception occurred as part of the execution. Examine the error message for details.

Examples – Configure or update integration configuration

```
$ vracli vrli set my-vrli.local
$ vracli vrli set 10.20.30.40

$ vracli vrli set --ca-file /etc/ssl/certs/ca.crt 10.20.30.40

$ vracli vrli set --ca-cert "$(cat /etc/ssl/certs/ca.crt)" 10.20.30.40

$ vracli vrli set --insecure http://my-vrli.local:8080

$ vracli vrli set --agent-id my-vrli-agent my-vrli.local

$ vracli vrli set --environment staging my-vrli.local
```

Clear integration of vRealize Log Insight

Command

```
vracli vrli unset
```

Arguments

There are no command line arguments.

Output

Confirmation is output in plain text format.

Exit codes

The following exit codes are possible:

- 0 - The configuration was cleared or no configuration existed.
- 1 - An exception occurred as part of the execution. Examine the error message for details.

Examples – Clear integration

```
$ vracli vrli unset  
Clearing vRLI integration configuration  
  
$ vracli vrli unset  
No vRLI integration configured
```

Participating in the Customer Experience Improvement Program for vRealize Automation

5

This product participates in VMware's Customer Experience Improvement Program (CEIP). The CEIP provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products.

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

This chapter includes the following topics:

- [How do I join or leave the Customer Experience Improvement Program for vRealize Automation](#)
- [How do I configure the data collection time for the Customer Experience Improvement Program for vRealize Automation](#)

How do I join or leave the Customer Experience Improvement Program for vRealize Automation

Join or leave the Customer Experience Improvement Program (CEIP) from the vRealize Automation appliance command line.

You can join the CEIP program when you install vRealize Automation and with the vRealize Lifecycle Manager (LCM). You can also join or leave the program by using command line options after installation.

To join the Customer Experience Improvement Program by using command line options:

- 1 Log in to the vRealize Automation appliance command line as **root**.
- 2 Run the `vracli ceip on` command.
- 3 Review the Customer Experience Improvement Program information and run the `vracli ceip on --acknowledge-ceip` command.
- 4 To restart the vRealize Automation services, run the `/opt/scripts/deploy.sh` command.

To leave the Customer Experience Improvement Program by using command line options:

- 1 Log in to the vRealize Automation appliance command line as **root**.

- 2 Run the `vracli ceip off` command.
- 3 To restart the vRealize Automation services, run the `/opt/scripts/deploy.sh` command.

How do I configure the data collection time for the Customer Experience Improvement Program for vRealize Automation

You can set the day and time when the Customer Experience Improvement Program (CEIP) sends data to VMware.

Procedure

- 1 Log in to the vRealize Automation appliance command line as **root**.
- 2 Open the following file in a text editor.
`/etc/telemetry/telemetry-collector-vami.properties`
- 3 Edit the properties for day of week (dow) and hour of day (hod).

Property	Description
<code>frequency.dow=<day-of-week></code>	Day when data collection occurs.
<code>frequency.hod=<hour-of-day></code>	Local time of day when data collection occurs. Possible values are 0–23.

- 4 Save and close `telemetry-collector-vami.properties`.
- 5 Apply the settings by entering the following command.

```
vcac-config telemetry-config-update --update-info
```

Changes are applied to all nodes in your deployment.